A model-based thruster leakage monitor for the Cassini spacecraft

Allan Y. Lee
M.S. 198-235
Jet Propulsion Laboratory
California Institute of Technology
4800 Oak Grove Drive, Pasadena, CA 91109-8099.

Abstract

The Cassini spacecraft was launched on October 15, 1997. It uses thrusters to perform many spacecraft control functions: to detumble the spacecraft, to maintain three-axis attitude control, to perform small trajectory correction burns, to desaturate the reaction wheels, and others. If one of the prime thrusters leaks (e.g., becomes stuck open), the opposing thrusters will be fired "extra hard" to maintain the commanded spacecraft pointing attitude. Obviously, the resultant draining of hydrazine and the excessive firings of thrusters cannot be allowed to persist indefinitely. A set of three model-based thruster leakage detection monitors has been designed, tested, and implemented in the flight software to protect Cassini against the occurrence of this highly unlikely event. Test results indicated that these monitors can detect small thruster leakage quickly and robustly. Once the leak is detected, the onboard fault protection logic will activate corrective actions, including the swapping of thruster branches.

1. Cassini and its mission

The Cassini spacecraft (S/C) was launched on October 15, 1997 by a Titan 4B launch vehicle. After an interplanetary cruise of almost seven years, it will arrive at Saturn in July 2004. To save propellant, Cassini will make several gravity-assist flybys: two at Venus, and one each at Earth and Jupiter. Fig. 1 depicts a pre-launch version of the interplanetary trajectory design of the Cassini spacecraft.

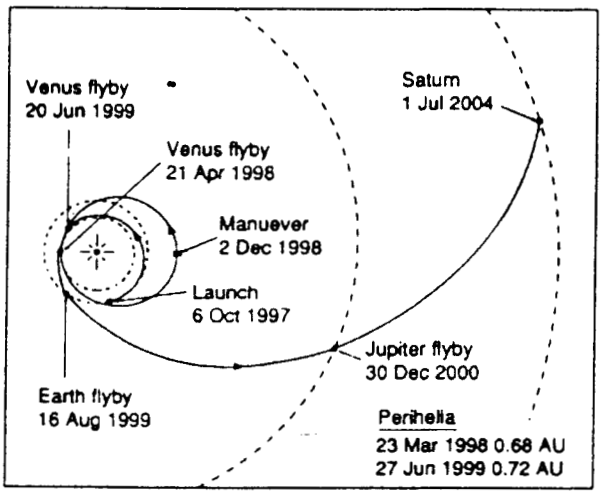

Fig. 1 Cassini interplanetary trajectory

Unlike both Voyager 1 and 2, which only flew by Saturn, Cassini will orbit the planet for at least four years. Major science objectives of the Cassini mission include investigations of: the configuration and dynamics of Saturn's magnetosphere, the structure and composition of the rings, the characterizations of several of Saturn's icy moons, and others. The Huygens probe, developed by the European Space Agency, will be released in September 2004, and will study the atmosphere of Titan, the only moon in the Solar system with a substantial atmosphere. Detailed descriptions of various science instruments carried onboard the Cassini spacecraft are given in

Ref. 1.

The Cassini spacecraft with its magnetometer boom deployed is depicted in Fig. 2.

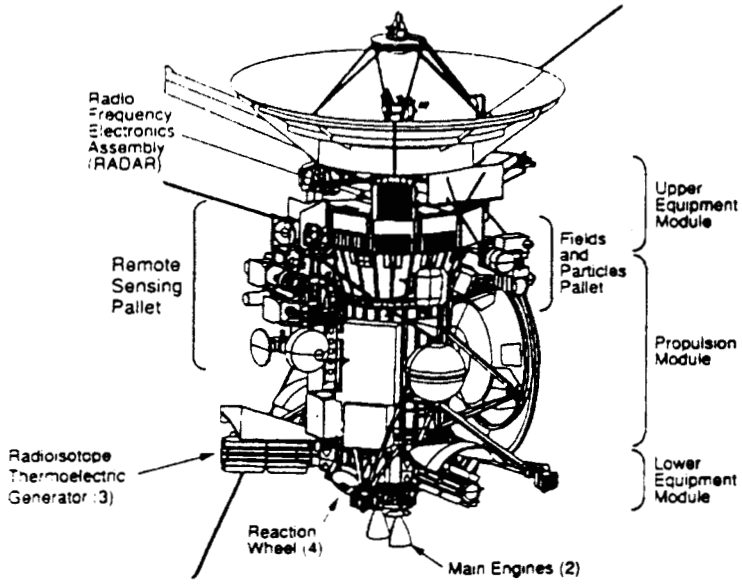

Fig. 2 The Cassini spacecraft

The central structure of the S/C carries the propulsion module with tanks for both the mono-propellant and bi-propellant. It also supports the Huygens probe and the remote sensing pallet (on which various high-resolution imaging cameras are mounted). Above the central structure is the upper equipment module and the high-gain antenna (HGA). One of the two low-gain antennas is collocated and co-boresighted with the HGA, while the second is mounted just underneath the probe. Below the central structure is the lower equipment module, which supports three radioisotope thermoelectric generators (power generators). Below the lower equipment module is a redundant set of two rocket engine assemblies. The lower equipment module also supports a structure on which are mounted four thruster pods.

At launch, the total mass of the S/C is 5573.8 kg and its inertia tensor is [8782 -124 114; -124 9050 -50; 114 -50 3773] kg-m^2. The knowledge uncertainties of the mass and moments of inertia of the "dry" S/C are less than $\pm 0.5\%$ and $\pm 10\%$, respectively, while the uncertainty of the products of inertias is less than ± 75 kg-m^2. The knowledge uncertainty of the S/C's center of mass is less than ± 5 cm. As mono-propellant and bi-propellant are being depleted over time due to thruster firings and main engine burns, there will be continuous changes in these S/C's inertia properties. Additionally, large changes in these inertia properties will occur at the following discrete events: deployment of the magnetometer boom, ejection of the Huygens probe, and after the long Saturn orbit injection burn. Knowledge uncertainties associated with these S/C inertia properties as well as those associated with the thrusters (to be described in Section 2) make the task of designing a thruster leakage detection scheme very challenging.

2. Thrusters and their functions

Cassini uses a set of eight thrusters to maintain three-axis attitude control of the spacecraft. Fig. 3 depicts the locations of the four thruster cluster pods.

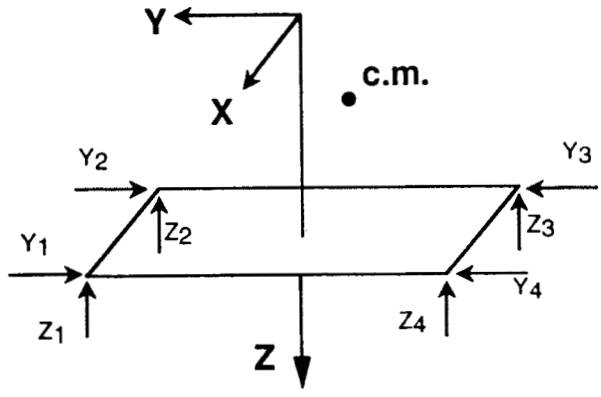

Fig. 3 Thruster pod locations

Pointing controls about the S/C's X and Y axes are performed using four Z-facing thrusters. Controls about the Z axis are performed using four Y-facing thrusters. The eight thrusters on the A-branch are being "backuped" by another set of eight thrusters on the B-branch. In addition to performing the vital task of pointing control, thrusters also serve the following functions: to detumble the S/C and bring it to a quiescent state after the Centaur/Cassini and Huygens/Cassini separations, to perform a full-sky search to locate the Sun, to slew the S/C about a selected

axis, to periodically unload the angular momenta that are accumulated on the reaction wheel assemblies (RWAs), and to perform small trajectory correction maneuvers (TCMs, which are sometimes called ΔV burns). Algorithms that perform these thruster-based spacecraft control functions are described in Ref. 2.

The monopropellant propulsion system for Cassini is of the blow-down type with one planned recharge. With this system, the hydrazine tank pressure, which is ≈2550 kPa at launch, will decay slowly with time as hydrazine is being depleted due to thruster firings. The nominal thrust of these thrusters at launch is ≈1 N. Knowledge uncertainties of the thrust produced by these thrusters include a ±7% pulse-to-pulse and a ±5% thruster-to-thruster variations.[3] The angular misalignment between the actual thrust vector and the S/C axes is below 1°, and the locations of the thruster nozzle centers are known to better than 1.1 cm.[3]

3. Leakage detection requirements

If one of the eight prime thrusters leaks (e.g., becomes stuck open), the expulsing hydrazine will impart angular momenta on two of the three S/C axes. In response to the resultant attitude control errors that appeared on the affected axes, appropriate thrusters will be fired to maintain the commanded spacecraft attitude. Obviously, the draining of hydrazine and the excessive firings of thrusters cannot be allowed to persist indefinitely. A set of three leak detection monitors (one for each S/C axis) has been designed, tested, and implemented in the flight software to protect Cassini against the occurrence of this highly unlikely event. Once a leak is detected, the onboard fault protection logic will activate autonomous corrective actions, including the swapping of thruster branches.

The main requirement on the thruster leakage detection monitor design is spelled out in Ref. 4: "... In the 'Cruise', 'Earth Approach', and 'Venus/Earth Flybys' modes, the Attitude Control Subsystem shall be able to detect any single thruster leak that applies an average torque of at least 0.005, 0.001, and 0.05 Nm, respectively, about any S/C axis, and shall isolate that single thruster leak before it applies more than 100 Nms of angular momentum about any S/C axis.···" In the Cruise mode, Cassini Mission Operations will communicate with the spacecraft at least once a week. Over this one-week time span, a thruster leak that resulted in a disturbance torque of 0.005 Nm will drain not more than 1.15 kg of hydrazine. This represents only 0.8% of the total hydrazine load. During the Venus 1, Venus 2, and Earth flybys, gravity gradient torques exerted on the S/C axes are comparable with 0.005 Nm. Hence, the torque detection level is relaxed (increased) to 0.05 Nm. On the other hand, during the 'Earth approach' phase of the mission, the spacecraft trajectory must be controlled very precisely. Therefore the torque detection requirement is tightened (decreased) to 0.001 Nm. Fig. 4 depicts the variation of the thruster leakage torque detection requirement over the first two years of the Cassini mission.

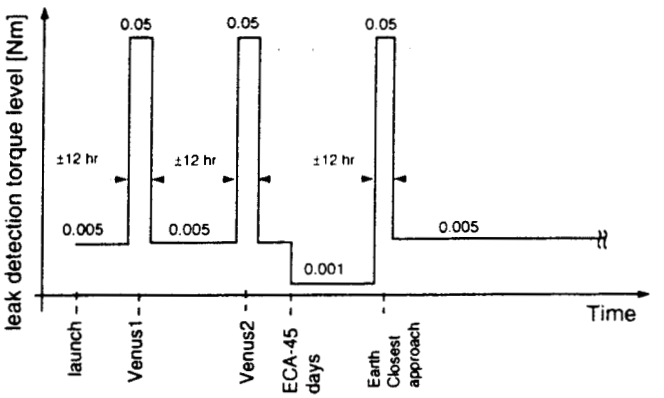

Fig. 4 Thruster leakage torque detection levels

Implicit in the above stated requirement is the need to detect thruster leaks even while the S/C is in a nonquiescent state (e.g., during a 'spiral' sun search maneuver). A thruster leakage detection error monitor design that meets this requirement will be described in Section 4. However, if instead the requirement was to detect thruster leaks only when the spacecraft is in a quiescent state, then the approaches described in the Appendix, which are simpler, could do the same job.

4. Leak detection monitor design

Conventional fault detection methods typically involve the monitoring of one or more of the following quantities: measured system output(s), estimated system state(s), and estimated process parameter(s).[5] These measured or estimated quantities are then compared with their "normal" values, and their deviations computed. If anyone of these deviations persistently exceeds its normal (allowable) tolerance, an error monitor is triggered to report this abnormality.

This conventional error detection approach is not applicable here because there isn't any single measured or estimated quantity that signals the presence of a leaky thruster. A different approach must be taken in designing the thruster leakage detection monitors. To this end, we first note that the rotational motion of the S/C is governed by the following Euler's equation:

$$I\dot{\vec{\omega}}+\vec{\omega}\times(I\vec{\omega}+\vec{H}_{\rm rwa}) = \vec{T}_{\rm rwa}+\vec{T}_{\rm pms}+\vec{T}_{\rm env}+\vec{T}_{\rm leak}+\vec{\epsilon}.$$

In this equation, I is the S/C inertia tensor. The S/C angular velocity vector, $\vec{\omega}$, is estimated by an onboard attitude estimator but it is typically noisy. Hence, the commanded angular velocity vector generated by an onboard attitude commander is used instead. The S/C angular acceleration vector, $\dot{\vec{\omega}}$, is not estimated by the attitude estimator. Again, we use that generated by the attitude commander. The angular momentum vector of the three RWAs, $\vec{H}_{\rm rwa}$, is added to the S/C angular momentum, $I\vec{\omega}$, to generate the total spacecraft angular momentum vector.

Reaction torques exerted on the S/C by the RWAs, $\vec{T}_{\rm rwa}$, is estimated onboard by a "RWA manager." Torque due to thruster firings, $\vec{T}_{\rm pms}$, is not available directly. Instead, an onboard "Propulsion manager" estimates the force impulses due to all prime thruster firings. Using the estimated thruster moment arms, these force impulses are next converted into three per-axis torque impulses. In effect, what we have is $\int \vec{T}_{\rm pms}\, dt$. These estimated angular momentum impulses are typically noisy, and contain errors due to knowledge uncertainties associated with the thrusters involved.

Environmental torques due to gravity gradient, solar radiation, magnetic field, atmospheric, etc. are captured in $\vec{T}_{\rm env}$. These torques are typically very small except during planet and Titan flybys. Torque due to a leaky thruster is represented by $\vec{T}_{\rm leak}$. For example, if the Z_1 thruster leaks, there will be a negative and a positive torque acting on the S/C X and Y axes, respectively (but with no torque on the Z axis). Finally, $\vec{\epsilon}$ is used to account for both the knowledge uncertainties associated with various S/C parameters (e.g., inertia tensor) and estimation errors associated with various derived S/C variables (e.g., angular momentum impulses).

Next, we note that various terms in the Euler's equation are available onboard. This allows us to monitor the entire Euler's equation instead of monitoring an individual measured/estimated quantity. Nominally, the Euler's equation does not contain the $\vec{T}_{\rm leak}$ term, and any deviation from this condition indicates a possible fault. Let us define the following residual angular momentum vector, $\vec{R}$:

$$\begin{aligned}\vec{R}(t) &= \int_0^t \{\, \vec{T}_{\rm leak} + \vec{\epsilon}\,\} d\tau \\ &= \int_0^t \{I\dot{\vec{\omega}} + \vec{\omega}\times(I\vec{\omega}+\vec{H}_{\rm rwa}) - \vec{T}_{\rm rwa} - \vec{T}^{f}_{\rm pms}\} d\tau.\end{aligned}$$

Here, $\vec{T}^{f}_{\rm pms}$ denotes a filtered version of the noisy raw data, and $\vec{T}_{\rm env}$ has been neglected in the last equation. With no leak, $\vec{R}$ contains only small zero-mean random fluctuations, due to $\vec{\epsilon}$. Whenever a leak appears, two of the three $\vec{T}_{\rm leak}$ components will be nonzero. If the leak persists, the resultant nonzero components of $\vec{R}(t)$ will grow (either increase or decrease) with time. In time, one of these components will exceed a pre-selected angular momentum threshold (R_T): $\mid R_i \mid \geq R_T$ (where i = x, y, or z axis), triggering the error monitor for that particular S/C axis. Note also that the polarities of R_i (i = x, y, and z axis) reveal the identity of the leaky thruster. See Table 1 for a complete mapping between a leaky thruster and the polarities of the resultant R_i (i = x, y, and z axis).

Table 1 Identity of the leaky thruster

Leaky thruster	Residual angular momenta [Nms]		
	R_x	R_y	R_z
Y_1	–	0	+
Y_2	–	0	–
Y_3	+	0	+
Y_4	+	0	–
Z_1	+	–	0
Z_2	+	+	0
Z_3	–	+	0
Z_4	–	–	0

In order to detect a leak before it imparts more than 100 Nms on any S/C axis, we select R_T to be 50 Nms. The rationale for this selection is as follows. The first time the 50-Nms threshold is exceeded, a fault protection activation rule, in attempting to stop the leak, will reset the controller unit for the thruster valve drive electronics. At the same time, all three components of $\vec{R}$ are reset to zero. If the reset of the controller unit does not stop the leak, and if the leak persisted, the same two $\vec{R}$ components will again grow with time. The second time the 50-Nms threshold is exceeded, fault protection activation rules will initiate swapping of thruster branches, which will stop the leak. In this way, we will be able to stop the leak before it imparts a total of 100 Nms on any spacecraft axis. The selected 50-Nms angular momentum threshold is changeable via a command.

5. Coping with uncertainties

Not all the uncertainty terms that affect $\vec{R}$ are random in nature. In particular, the thruster-to-thruster variation can impart systematic errors to the Euler's equation, causing $\vec{R}$ to grow with each thruster firing, even without a leak. As such, prolonged thruster firings could trigger the $\mid R_i \mid \geq R_T$ criterion. Two modifications are made to the above described leak detection scheme to avoid such a false alarm.

First, a correction vector $\vec{R}^{\text{corr}}$ is estimated to account for angular momentum accumulation due to thruster-to-thruster variation. The estimation of the X-axis component of $\vec{R}^{\text{corr}}$ is given here as an illustration.

Rotations about the spacecraft X axis are made using either the Z_1 and Z_2 pair or the Z_3 and Z_4 pair. Hence, the X-axis component of $\vec{R}^{\text{corr}}$ is proportional to both the differential thrusters' impulses and the size of the thruster-to-thruster uncertainty estimate:

$$\begin{aligned} \Delta I_X &= \mid \Delta I_{Z_3} + \Delta I_{Z_4} - \Delta I_{Z_1} - \Delta I_{Z_2} \mid, \\ R_X^{\text{corr}} &= 2 \times \eta \times 1.6 \times \Delta I_X \,. \end{aligned}$$

Here, ΔI_{Z_j} is the accumulated impulses due to the Z_j thruster over the last RTI (real time interrupt). They are available from the on-board "Propulsion" manager. The thruster-to-thruster uncertainty η is estimated to be 0.05 (i.e., 5%) at launch. If a better estimate of η is available from post-launch telemetry data, the current estimate is changeable using the same command that is used to alter "R_T". The factor "1.6" represents the average magnitude of the four Z thrusters' moment arms (in meters). The factor "2" is used to account for the fact that the standard deviation of a variable that is the sum/difference of four variables, each with a standard deviation of σ, is 2σ. The Y and Z components of $\vec{R}^{\text{corr}}$ are estimated similarly. The correction vector $\vec{R}^{\text{corr}}$ is then added to $\vec{R}$, and the resultant modified $\vec{R}$ is used in the momentum threshold check.

Next, after the angular momentum threshold is exceeded, a second condition: $T_j^{\text{trigger}} \leq T^{\text{limit}}$ is checked before any fault protection action is initiated. Here, T_j^{trigger} (j is the S/C axis whose $\mid R_j \mid$ exceeded the angular momentum threshold) is the time it takes $\mid R_j \mid$ to exceed R_T since it was last reset. The threshold T^{limit} is a "time-domain" threshold that is to be pre-selected. It is changeable using the same command that is used to alter "R_T". If T_j^{trigger} is larger than T^{limit}, then we believe that the R_T threshold was exceeded not because of a leak but is due rather to the systematic accumulation of angular momentum from prolonged thruster firings. On the other hand, if T_j^{trigger} is smaller than T^{limit}, then there is a "real" leak, and corrective action from the onboard fault protection logic is needed. This time-domain trigger criterion is illustrated in Fig. 5.

For the 'Cruise' mode, we select T^{limit} to be 10 hours. Note that this value is larger than 50/0.005/3600≈ 2.8 hours. Hence, a leak detection level of 0.005 Nm could be met. Early Cassini flight data indicate that hydrazine is being consumed at a rate of about 1-1.5 grams/day. This consumption rate is likely to decrease once the spacecraft gets further away from the Sun. Using a worst-case hydrazine consumption rate of 1.5 g/day, an upper bound on the per-axis angular momentum (due to thruster uncertainty) that is accumulated over 10 hours is ≈0.05 Nms. It is about three orders of magnitude smaller than R_T (50 Nms). As such, T^{limit}= 10 hours is a good choice for the 'Cruise' phase of the Cassini mission. The time threshold T^{limit} is to be changed several times througout the mission to reflect the changing torque detection levels described in Section 3 (see also Fig. 4).

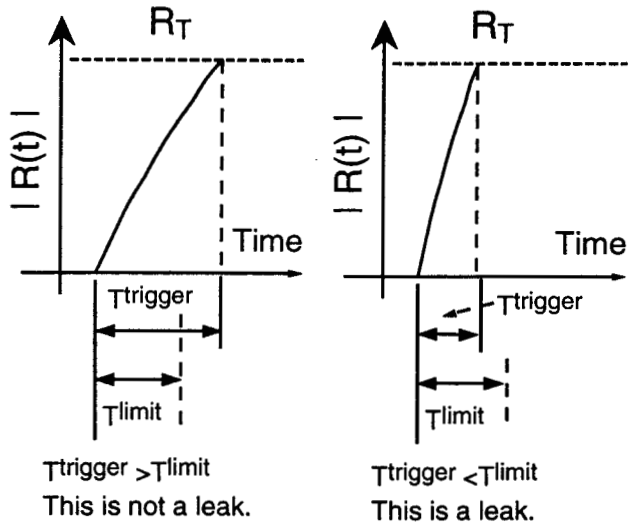

Fig. 5 Leak detection scheme

6. Simulation results

The coded leak detection design was tested using the Cassini Flight Software Development Testbed. Attitude control actuators (e.g., thrusters), attitude determination sensors, and the spacecraft itself are represented by validated analytical models in this testbed. The testbed can simulate both the pulse-to-pulse and thruster-to-thruster variations, as well as disturbance torques due to a leaky thruster.

The leak detection monitor is designed to perform its function in various spacecraft scenarios: during a spiral sun search, while the S/C is being slewed from one target to another, and others. As such, its performance in a large combinations of scenarios have to be verified: (a) either a leaky Z or Y thruster, (b) either the leaky thruster is on the A or B thruster branch, (c) either the leak is small or large (from 0.1% to 100%), and (d) either the leak occurred when the spacecraft is in a quiescent state, in a constant rate slew, or is being accelerated/decelearted about a spacecraft (either X, Y, or Z) axis. For brevity, only results obtained for a scenario in which the Z_1 thruster develops a 10% leak while the S/C is being slewed about the Y axis are given here.

A 10% leak in the Z_1 thruster will generate disturbance torques of –0.158 and +0.124 Nm about the S/C X and Y axes, respectively. These torques will cause the X and Y components of $\vec{R}$ to grow with time (cf. Fig. 6). About 6 minutes (which is well below T^{limit}) latter, | R_x | first exceeded R_T, causing the X-axis error monitor to be triggered. In this particular simulation, we assume that the first corrective action initiated by the fault protection logic (the reseting of the thruster controller unit) does not stop the leak (as depicted in Fig. 6). Accordingly, both R_x and R_y will continue to grow with time after having been reset to zero. The next triggering of the error monitor will lead to the swapping of thruster branches, which stops the leak.

7. Summary

A set of three model-based thruster leakage detection monitors has been designed, tested, and implemented in the flight software to protect Cassini against the occurrence of a thruster leak, a highly unlikely event. The computational effort involved in executing this set of monitors by the flight computer is moderate. With only two thresholds to select, this set of monitors could be easily managed by Mission Operations controllers. Simulation test results indicated that the design meets all the requirements stated in Section 3. In particular, the design can detect thruster leaks that are ≈0.1% of the thruster

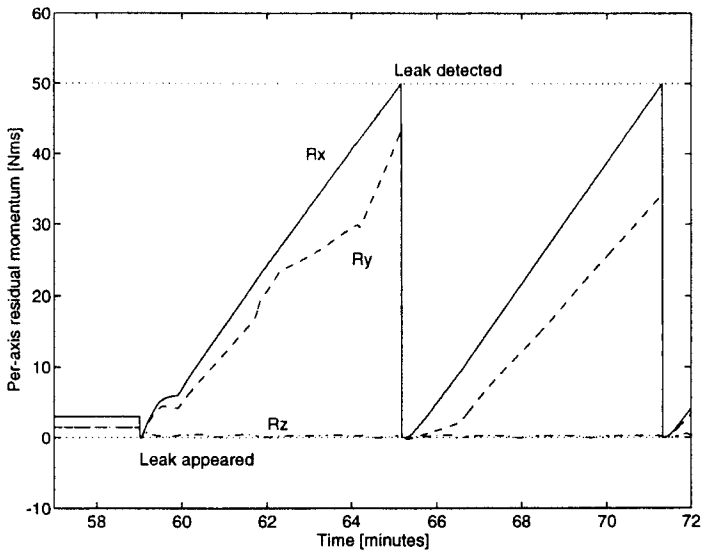

Fig. 6 Time histories of residual momenta

magnitude, and does so quickly before an unacceptable level of angular momentum is imparted on any spacecraft axis. The robustness of the design against knowledge uncertainties of various spacecraft parameters as well as estimation erros of various spacecraft derived variables has also been confirmed via extensive simulations.

8. References

1. Jaffe, L. and Herrell, L., Cassini Huygens Science Instruments, Spacecraft, and Mission, Journal of Spacecraft and Rockets, Vol. 34, No. 4, July-August, 1997.

2. Wong, E., and Breckenridge, W., An Attitude Control Design for the Cassini Spacecraft," Proceedings, AIAA Guidance, Navigation, and Control Conference, August 1995.

3. Lee, A., Cassini Orbiter Functional Requirements Book: Accuracy Requirements and System Capabilities, CAS-3-170, JPL D-699-205, Revision D, November 1997.

4. Walker, J., Cassini Orbiter Functional Requirements Book: Attitude and Articulation Control Subsystem, CAS-4-2007, JPL D-699-205, Revision F, November 1997.

5. Isermann, R., "Process Fault Detection Based on Modeling and Estimation Methods - A Survey," *Automatica*, Vol. 20, No. 4, pps. 387-404, 1984.

9. Acknowledgments

The research described in this paper was carried out by the Jet Propulsion Laboratory, California Institute of Technology, and was sponsored by the National Aeronautics and Space Administration. The author wishes to thank T. Barber, W. Breckenridge, G. M. Brown, M. J. Brown, J. Chodas, C. Jennings, S. Johnson, M. Lam, M. Leeds, G. Macala, A. Mark, S. Peer, R. Rasmussen, and D. Skulsky, his colleagues at JPL, and K. Hilbert, formerly with JPL, for many helpful discussions. The methodology described in this paper has been reported to the Patent Counsel Office of Jet Propulsion Laboratory on June, 1998 for potential patent application.

Appendix

Which thruster leaks?

One source of telemetry data that could be used to identify the leaky thruster is the time histories of the per-axis attitude control errors (which are denoted by e_i, i = x, y, and z axis). As indicated in Table A1, since a leaky Y_1 thruster generates both a positive and a negative torque about the X and Z axes, respectively, e_x and e_z will be "stuck" at $-db_x$ and $+db_z$, respectively (db_i = attitude controller deadband for the S/C's i^{th} axis). With no torque on the Y axis, e_y will fluctuate between $\pm db_y$. By studying the time histories of these e_i components, we can both detect the presence of a leaky thruster and identify which thruster leaks. Attitude controller error signatures of other leaky thrusters are tabulated in Table A1.

Table A1 Per-axis attitude error signatures

Leaky thruster	Per-axis attitude error [mrad]		
	e_x	e_y	e_z
Y_1	$-db_x$	fluctuate	$+db_z$
Y_2	$-db_x$	fluctuate	$-db_z$
Y_3	$+db_x$	fluctuate	$+db_z$
Y_4	$+db_x$	fluctuate	$-db_z$
Z_1	$+db_x$	$-db_y$	fluctuate
Z_2	$+db_x$	$+db_y$	fluctuate
Z_3	$-db_x$	$+db_y$	fluctuate
Z_4	$-db_x$	$-db_y$	fluctuate

Another source of telemetry data that could be used to detect the presence of a leaky thruster is the accumulated ontime's of the eight prime thrusters (T_j^{on}, where j = Y_1-Y_4 and Z_1-Z_4). Specifically, the time rates of increase of these ontimes, $\dot{T}_j^{on}$, will reveal the identity of the leaking thruster.

Consider the scenario when the Z_1 thruster leaks. A negative and positive external torque will soon appear on the S/C's X and Y axes, respectively. With this combination of external torques, the Z_1 thruster, which is leaking, will not be fired. On the other hand, the Z_3 thruster, which is located diagonally opposite the leaky Z_1 thruster, will be fired at a rate that is about double that of the other two Z thrusters (Z_2 and Z_4). This is because the Z_3 thruster must negate external torques that appeared on both the S/C's X and Y axes. Accordingly, we place "r" in both the Z_2 and Z_4 columns of the Z_1 row in Table A2, while that for the Z_3 column is "2r". Without any external torque on the S/C's Z-axis, none of the Y thrusters will be fired. Hence, the rates of increase of all the Y thrusters' ontimes are negligible (accordingly, we place "0" in the Y_1-Y_4 columns of the Z_1 row in Table A2). A complete mapping between the leaky thruster and the thrusters' ontime rates is given in Table A2.

Table A2 Signatures of $\dot{T}^{on}$

Leaky thruster	Thruster ontime rates [sec/s]							
	Y_1	Y_2	Y_3	Y_4	Z_1	Z_2	Z_3	Z_4
Y_1	0	r	0	r	r	r	0	0
Y_2	r	0	r	0	r	r	0	0
Y_3	0	r	0	r	0	0	r	r
Y_4	r	0	r	0	0	0	r	r
Z_1	0	0	0	0	0	r	2r	r
Z_2	0	0	0	0	r	0	r	2r
Z_3	0	0	0	0	2r	r	0	r
Z_4	0	0	0	0	r	2r	r	0